DOI:10.1145/3786780

Steffen Wendzel, Julia Lenz, and Martin Baumann

Opinion

Internet Censorship's Implications on Academic Research: Motivating an Academic Accessibility by Design

Considering the influence of Internet censorship on academic research and motivating an accessibility by design for research artifacts.

Internet-based fake news is often considered as the main problem when it comes to global dissemination of information. Although the fake news issue is indeed of utmost importance, as it introduces false information into public communication channels, there is also another major problem that can be considered its sibling: the prevention of access to information through Internet censorship.

During the past decade, censorship has been examined from many angles. Most of them analyzed the censorship methodology, conducted measurements of Internet censorship, improved the methodology of such measurements, or focused on improving or evaluating censorship circumvention techniques.[4,7,10]

An important societal role of Internet censorship research is that it helps to reveal when and where governments interfere with online content. These insights are especially valuable around elections or other high-stakes political events. Moreover, evaluating whether Internet censorship has been conducted can become important long after censoring events took place. Although the political level has been subject to several studies, a neglected aspect of Internet censorship is its effect on the global academic freedom and exchange. We would like to emphasize that no country can be considered entirely immune to limitations of academic freedom, at least not in the long run.

IMAGE BY GANNVECTOR

Inaccessible Academic Work

The most obvious problem is limited or blocked access to academic publications. Typically, well-funded institutions can obtain paywalled papers, but censors may still block publications discussing "undesired" topics. For instance, a researcher (or even an interested citizen) who tries to retrieve studies discussing minority rights or local censorship behavior of the country they are residing in, may find the pages unreachable. Keyword filters or

agreements with publishers can simply erase those articles from view.

However, we should not neglect the fact that scientific artifacts are more than journal articles. Datasets, code, software repositories—sometimes even the very tools used to circumvent censorship, like Tor—are equally vulnerable or become inaccessible because the binaries or the software code are hosted on a blocked domain.

GFWatch,[5] for example, is a project that monitors China's Great Firewall (GFW) filtering system. The GFW is one of the most sophisticated filter systems on the planet, and—thanks to GFWatch—domain filter rules of the GFW can be extracted. When we analyze the extracted filter heuristics, we can spot github.io subdomains and several university websites (for example, London School of Economics; lse.ac.uk). Moreover, several important websites that are linked to the provision of open source software are blocked, for example, redhat.com. While there are only few such results today, and some could even be the result of collateral damage, they cover solely China's filter system and there is no guarantee that the number of these results will remain as is. Similar sophistication of filter systems can be found in other countries, too. For instance, Iran aims to implement a "Halal Internet" with restrictive filter technology. Moreover, temporary blockings of selected domains are not easy to detect. Several countries have already been reported to carry out complete Internet shutdowns, affecting not only social media sites but also academic freedom. However, according to the latest Freedom on the Net (FOTN) report,[4] Internet censorship is not limited to authoritarian countries. Instead, several democratic countries are threatened by it as well. The report highlights that the Internet's freedom recently has decreased in 27 of the 72 covered countries while it improved in only 17, and that during the last 15 years of analysis, "state authorities and other powerful actors [...] transformed the ways [in which they] assert control over information."[4] With the increasing attempts to censor the Internet, it is also becoming more likely that some researchers around the globe are not able to access a certain publication, dataset, or software at a given moment in time.

No country can be considered entirely immune to limitations of academic freedom, at least not in the long run.

Moreover, AI-driven systems, such as chatbots, that support academic research are increasingly gaining popularity. Both their training and inference phase as well as the utilization of such tools might be influenced by censorship. Several implications arise for research: (1) If certain sensitive material is blocked during data collection and training, the model cannot reproduce facts it never learned, creating responses with censor-moderated bias.[3,4] Thus, resulting in limiting the insights provided to academics and leaving knowledge gaps where scholars need detail, such as minority health datasets or protest archives. Recent work showed that systematic removal of politically sensitive Chinese webpages during scraping leads to holes in factual recall and topic coverage.[1] (2) Language bias compounds the problem: If training data in a given language is censored, the model may refuse, omit, or distort information when prompted in that same language while providing fuller answers in less restricted languages.[8] These language-conditioned blind spots can encourage isolated "knowledge bubbles," limiting cross-disciplinary and cross-cultural exchange, mirroring the echo chambers already documented in social media. (3) Since large language models are increasingly integrated into research tools and can be expected to play a key role in future library catalogs and academic databases, even subtly incorrect responses, for example, as in the case of transphobia,[9] have the potential to influence and thus corrupt research results.

Human Factors: Research Not Being Conducted or Published

The mentioned reasons for inaccessibility of academic work do more than potentially twist content; they quietly convey what is "safe" to ask. When searches for certain topics contain errors or silence, scholars may quietly learn to sidestep those topics long before any formal ban. This makes the human dimension impossible to ignore. While we summarized several human aspects of Internet censorship in previous work,[10] we would like to emphasize how the mentioned technical, legal, social, and moral regulations translate those pressures into behavior. Technical censorship flows seamlessly into self-censorship: prompting academics (like all other citizens) to adjust their own behavior to avoid violating "rules." For instance, a researcher refrains from accessing a website hosting a particular paper repository if they suspect it might be a website considered unacceptable by the censor. Self-censorship starts with an internal cost-benefit calculation by asking oneself questions such as "Will this keyword flag a network monitor?"—and potentially end with proactive, selective avoidance. Another potential result of self-censorship is that academics might be afraid to publish results and give talks that could harm their (regional) standing, academic position, or personal safety. Risk assessment might become a routine part of the workflow, especially concerning decisions about what to publish, where to publish, and who to engage with. Non-given talks also influence the chance for building scientific networks and for the "third-mission" (outreach to the public).

At the heart of this effect lies fear that one's research results could be deemed controversial by governments or that collaboration with peers in certain countries might draw scrutiny. Over time, these concerns can become internalized to the point that a researcher possibly alters or abandons their work on a specific topic entirely. Shaped by perceived threats, this self-limiting mindset diminishes the richness and rigor of research.

Meanwhile, when many scholars engage in self-censorship, a chilling

effect emerges, that spreads through the broader academic community (worldwide). Historically associated with protecting First Amendment rights in debate over safeguarding free speech in the U.S., the chilling effect describes a climate of uncertainty that discourages collective expression even when no explicit ban exists. The term "chilling" implies a slowing, suppressing, discouraging or delaying. This effect could be expressed in hesitations by grant officers to fund politically sensitive research projects or journal editors to anticipate controversy and desk-reject certain manuscripts. At first, in some situations such choices may seem sensible, yet together they slow the production of knowledge or result in a decline of discourse, especially when unsubstantiated.

The extended chilling effect even goes one step further: it lingers after the immediate perceived risk has faded and migrates to settings where that risk never existed. Habits formed under (perceived) restriction, such as avoiding certain censored citations or keywords, can become so ingrained that academics carry them into new institutions or countries. Over time, these behaviors harden into participation and diversity gaps within the research landscape. Meanwhile, talent migration—students and researchers leaving censored environments for less-restricted ones—creates a twofold problem. On the one hand, it drains home institutions of expertise, slowing their scientific development ("brain drain"). On the other hand, those who relocate may have internalized caution so that they carry censorship "across borders," extending the chilling effect.

A related issue is the rise of age-verification laws. Though meant to protect minors, these systems often overclassify legitimate scientific or health material as "restricted," and even under frameworks such as the EU's Digital Service Act can cut young people off from essential information. This creates barriers for emerging researchers.

All these scenarios and consequences of the chilling effect lead to unpublished and less-propagated research and an overall asymmetric distribution of knowledge production. Whereas researchers in environments considered relatively "free" by the FOTN report[4] can more easily explore frontier topics, their peers in censored environments—and anyone collaborating with them—face legal, professional or personal barriers. This silence is costly: researchers in censored countries possess the most direct evidence of how those systems work and could tell us a lot about essential insights into, for example, censorship practices taking place in their countries.

Academic Accessibility by Design

Given the aforementioned issues, Internet censorship's influences on academic freedom cannot be neglected. The research community must come up with remedies for these problems, considering the heterogeneous environments in which scholars reside. We propose three joint solutions (each benefiting from the others) that can be expected to aid scientific research when confronted with restricted Internet access.

Censorship-resistant scientific repositories. One such case could be the provision of multiple mirrors for downloading software or datasets. These mirrors could be "shipped" in different ways, for example, as links in papers, as part of a documentation that is accessible to a researcher, or even through canonical DNS records (CNAME). Such mirrors for papers, code and data outside of classical websites and web-based databases could also aid accessibility. An example is the utilization of BitTorrent downloads or the propagation of a Signal or Telegram profile allowing the request of academic artifacts. Moreover, research has shown that some forms of content provision force a censor to accept collateral damage: content-delivery networks (CDNs) provide a large spectrum of web content (another

Internet censorship's influences on academic freedom cannot be neglected.

case are co-located virtual hosts, but these are easier for a censor to analyze through TLS' Server Name Indication (SNI[7]). If the IP ranges of CDNs are blocked because of some specific web-content, the censor is essentially forced to also block the remaining content of the CDN, resulting in undesired side effects.[7] These ideas are not exhaustive and all of them are linked to individual limitations. However, we believe that both, the computer science and digital library communities must come up with more (and better) ideas to ultimately end up with increased academic accessibility that comes "by design."

Saving (last) copies. In university and state libraries, it is already common practice to track last copies of publications before parceling out.[6] Recently, this approach has been transformed into the world of scientific online publications and is usually handled by librarians. We suggest that scientists take responsibility and proactively initiate the duplication of records into the aforementioned censorship-resistant scientific repositories. Such a duplication should be triggered in cases where anticipated political changes threaten the existence of scientific publications and data.

Post-publication de-anonymization. Further, an option to publish papers anonymously allows researchers to contribute their knowledge to science, despite being located in a restricted or censored environment. This might be especially relevant in cases where research activities can be conducted even though the research field is subject to censorship (for example, when no governmental funding is required). During anonymized authorship, the impact on personal safety is minimized. A better approach would be a pseudonymized one that allows a later de-pseudonymization, but solely by the researchers themselves. Such attempts were already made centuries ago. For instance, Lewis Carroll is actually a pseudonym of Charles L. Dodgson. Pseudonyms were also used in academic writing, even for whole groups, such as Nicolas Bourbaki—a pseudonym of a group of French mathematicians. Pseudonymization does not conflict with rigorous academic reviews and can still allow a consideration of ethical aspects.

A censorship-resistant provision of scientific artifacts is of utmost importance.

Using pseudonymization, and if supported by publishers and publication databases, a de-pseudonymization would enable the later merging of real and pseudonymized academic profiles in the relevant publication databases, such as ORCID, dblp, or publisher-specific ones (ACM DL, IEEE Xplore, and so forth). This provides scholars with the option to attract citations even under restricted environments, at least with a delay. Situations justifying a later de-pseudonymization could be very different ones: the researcher could move to a non- or less-restricted country. Alternatively, a restrictive (regional) regime might disappear after some years.

However, it would be helpful if authors indicate that (some of) the author names are in fact represented by pseudonyms as this has implications for copyright transfers. Law experts would need to clarify how the copyright transfer process could be split into a pseudonymized and post-pseudonymized time-span and how mixed authorship (true names plus pseudonyms on the same paper) can be accomplished.[a]

Finally, these solutions are not sufficient to address the entirety of relevant issues posed by Internet censorship. For instance, as pointed out by the anonymous reviewers, protecting scientists' search and research histories remains essential, as AI-driven surveillance increasingly monitors web-based tracking of scholars and platform activity. Knowing that queries or interactions may be logged can deter scholars from sensitive topics, potentially further amplifying the chilling effect.

a Papers with anonymous co-authors are already present in today's research. An example can be found in Anonymous et al.[2]

Conclusion

We observe an increasing trend in many different countries to control access to information on the Internet by strengthening their endeavors to conduct Internet censorship with continuous increasingly sophisticated censoring techniques. We explained that Internet censorship threatens the core of our scientific freedom and integrity. It has the potential to impede scientific progress as well as the digital sovereignty of researchers. For this reason, a censorship-resistant provision of scientific artifacts is of utmost importance. We propose to develop practical methods to achieve accessibility of academic work by design. Finally, becoming resilient against Internet censorship is a necessity that must be addressed by multiple stakeholders: the research community, librarians, and publishers. C

References

1. Ahmed, M. and Knockel, J. The impact of online censorship on LLMs. In *Proceedings of FOCI* (2024).
2. Anonymous et al. Triplet censors: Demystifying great Firewall's DNS censorship behavior. In *Proceedings of FOCI* (2020).
3. Berger, H. and Shavitt, Y. Measuring DNS censorship of generative AI platforms (2024).
4. Freedom House. *Freedom of the Net Report.* (2025); https://freedomhouse.org/sites/default/files/2025-11/Freedom_on_the_Net_2025_Digital.pdf.
5. GFWatch project website (2025); https://gfwatch.org/
6. Hassel, L. bwLastCopies—Identifying the last copy: Activities in the German state of Baden-Württemberg. *Library Management 44*, (2022), 361–366.
7. Master, A. and Garman, C. A worldwide view of nation-state Internet censorship. In *Proceedings of FOCI* (2023).
8. Noels, S. et al. What large language models do not talk about: An empirical study of moderation and censorship practices. In *Proceedings of ECML PKDD* (2025); 10.48550/arXiv.2504.03803
9. Scheuerman, M.K. et al. Transphobia is in the eye of the prompter: Trans-centered perspectives on large language models. *ACM Trans. Computer-Human Interaction. 32*, 5 (2025).
10. Wendzel, S. et al. A survey of Internet censorship and its measurement: Methodology, trends, and challenges. *Computers & Security* (2025); 10.1016/j.cose.2025.104732

Steffen Wendzel (steffen.wendzel@uni-ulm.de) is a professor at Ulm University, Germany, where he is the chair of the Institute of Information Resource Management (IRM) and the director of Ulm University's library and computing center ("kiz").

Julia Lenz (julia.lenz@uni-ulm.de) is a Ph.D. student at the Institute of Information Resource Management (IRM), Ulm University, Germany.

Martin Baumann (martin.baumann@uni-ulm.de) is a professor for human factors at the Institute of Psychology and Education at Ulm University, Germany.

The authors express their gratitude for the valuable feedback received from the anonymous reviewers.